A Model of Homogenous Input Demand Under Price Uncertainty

Frank A Wolak
Charles D. Kolstad

BEBR

A Model of Homogenous Input Demand Under Price Uncertainty

Frank A. Wolak
Stanford University

Charles D. Kolstad, Assistant Professor
Department of Economics

Abstract

Motivated by the common observation that firms simultaneously purchase a homogenous factor of production from a variety of suppliers each charging a different price, we derive and test the empirical validity of a model of homogeous input demand under price uncertainty which suggests that a firm trades off the average magnitude of input costs against its variability (risk) in selecting the optimal input supplier mix. Using recent work in time series econometrics, we suggest a general methodology for modeling the stochastic process generating the vector of random input prices. Our model of input choice uses the parameters of this stochastic process to compute the optimal supplier share for each time period. Our model of input choice implies a point estimate of the firm's marginal rate of substitution between risk and expected cost. The statistical significance from zero of this marginal rate of substitution provides a test of the empirical validity of the risk diversification model of input choice. From this magnitude we can compute an estimate of the amount a firm would pay above the current expected market price for supplies of this input with no price risk. In addition, our model yields estimates of the supplier specific risk characteristics of each input price series similar to the β coefficient in the Capital Asset Pricing Model. Our specific application is to the Japanese steam coal import market. Five suppliers compete in this market: China, USSR, South Africa, U.S. and Australia. There is considerable anecdotal evidence for the validity of our model as a description of this market. We find that our model is able to provide an economically plausible justification for three empirical anomalies unexplainable by a model based on the least expected cost criterion for input choice.

A Model of Homogeneous Input Demand Under Price Uncertainty

1. Introduction

The purpose of this paper is to derive and test the empirical validity of a model of homogenous input demand under price uncertainty. The motivation for this investigation is the common observation that firms simultaneously purchase a homogeneous factor of production from a variety of suppliers each charging a different price. In addition, there are many instances when the price from one supplier is consistently above that of all other suppliers for an extended period of time yet firms continue to purchase from this supplier. This observation seems to violate the expected cost minimization input choice criterion.[1] An attempt to explain these anomalies has suggested that firms trade off the level of expected input cost against its variability in deciding how to allocate total input demand across available suppliers. By purchasing inputs from a variety of suppliers the firm is diversifying away some of the price risk associated with satisfying demand from the single least expected cost supplier.[2]

Although the marginal rate of substitution (MRS) between risk and cost is not directly observable, we develop a methodology for empirically estimating its magnitude from a time-series of input purchases. By examining the statistical significance of the sign of this marginal rate of substitution, a test of the risk diversification hypothesis is possible. The MRS is a point estimate of the firm's risk preferences; from this an input price risk premium, the percentage above the current expected market price a firm would pay for riskless input supply, can be

---

[1] Assume, for the sake of simplicity, the price series are independent and identically distributed draws from a multivariate distribution. The null hypothesis of equal means for the prices becomes less likely the greater the number of observations that one price series remains above the others. Clearly, if firms were expected cost minimizers they would purchase all of this input from the least expected price supplier. Hence, in this simple case, the nonzero market share of the consistently high priced supplier is with high probability a violation of the expected cost minimization criterion of input choice.

[2] Previous authors (Batra and Ullah (1974), Sandmo (1971), and Blair (1974)) have theoretically examined the comparative statitics of firm behavior under input and output price uncertainty.

calculated. This risk diversification model of input demand also provides an equilibrium framework for discussing the relative risk characteristics of input suppliers and the effect of these risk characteristics on observed input choice.

For an empirical implementation of this model we have chosen the Japanese steam coal import market. This coal is primarily used in Japanese cement manufacturing and electricity generation facilities. Although this coal is supplied to a variety of different consumers in Japan, the Ministry of International Trade and Industry (MITI) is the centralized decisionmaker which co-ordinates all international steam coal transactions and hence is analogous to the firm in our model of input demand. We estimate and test the model of risk diversification as an explanation for the observed patterns of Japanese steam coal imports. This market provides an especially challenging test of the risk diversification hypothesis because of three puzzling observations on the time series properties of the vector of prices of steam coal in Yen from the available suppliers. The first puzzle is that the price of United States coal is above that of all other suppliers throughout the entire sample period, yet the United States supplies an average of 12 percent of all steam coal imported to Japan during this period. The second is that the price of steam coal from the Soviet Union is consistently below the price of all other suppliers throughout the sample although it consistently has the smallest share of the Japanese steam coal import market. The final puzzle is that South Africa and Australia have almost exactly the same mean price over the sample although for all observations over this same period the share of Japanese steam coal imports from Australia is consistently more than double that from South Africa. We find that the risk diversification model and the apparent risk characteristics that it implies for each supplier explains these empirical observations about the Japanese steam coal import market.

The remainder of this paper proceeds as follows. The next section introduces our notation and then derives our risk diversification model of input demand. Section 3 discusses the

econometric framework underlying the estimation of this model. This section treats the specification of a stochastic process describing the behavior of the vector of input prices over time and also describes the form and sources of other uncertainty in the model. Section 4 provides a brief overview of the market for Japanese steam coal, to match up the theoretical model of section 2 with the actual workings of this market. Section 5 will describe the application of this framework to the Japanese steam coal import market and the implementation of a test of the risk diversification hypothesis. In section 6 we present the general implications of modeling input demand under uncertainty within this risk diversification framework. For example, we are able to calculate the risk premium described earlier and a measure of market specific risk associated with each of the supply price processes. Subject to the restriction that on the average the market is in equilibrium we can also derive a relationship between these measures of market specific risk and the expected supply price for each supplier. The paper closes with a short discussion of the policy implications of the empirical results and suggestions for future applications of this framework.

2. A Risk Diversification Model of Firm Input Demand

Consider a firm using a set of inputs to produce one or more outputs. All inputs to production but one are termed "non-risky" in that their price is non-stochastic. Output prices are also non-stochastic. The price of one of the inputs (the "risky" input) is uncertain. Supplies of that input must be contracted for *ex ante* before the price uncertainty is resolved. This is consistent with conventions used in other models of input choice under price uncertainty (e.g., Batra and Ullah (1979)). If the firm is risk averse, it may increase its utility by substituting away from the risky input or by utilizing a variety of suppliers in an effort to reduce risk through diversification.

We begin by defining our notation.

p_{it} price of risky input from supplier i in period t (i=1,...,n)

q_{it} quantity of risky input received from supplier i in period t (i=1,...,n)

p_t n-dimensional vector of risky input prices in period t

q_t n-dimensional vector of risky input quantities in period t

r_t vector of prices of non-risky inputs in period t

s_t vector of quantities of non-risky inputs in period t

π_t vector of deterministic output prices in period t

y_t vector of output quantities in period t

I_t information set available to firm at time t containing p_s, $(s \leq t-1)$

μ_t $E(p_t \mid I_t)$, conditional expectation of p_t

Σ_t $E\{(p_t - \mu_t)(p_t - \mu_t)' \mid I_t\}$, conditional variance of p_t

ι n-dimensional vector of 1's

Q_t $\iota'q_t$, total demand for risky input in period t

w_t q_t/Q_t, n-dimensional vector of risky input quantity shares.

The firm is governed by the implicit production relation $f(y_t,s_t,Q_t) = 0$. Rather than maximize profits, because it is risk averse, in each period the firm maximizes the expected utility of profits given the vectors of nonstochastic input and output prices and the information set I_t. We make the simplifying assumption that the firm's expected utility can be written as a function of the conditional expectation of profits, $E(\Pi_t \mid I_t)$, and the conditional variance of profits, $V(\Pi_t \mid I_t)$, where

$$\Pi_t = \pi_t'y_t - p_t'q_t - r_t's_t$$

is the firm's profit in period t. This assumption about firm preferences is similar to that made for investor preferences in the Capital Asset Pricing Model (CAPM). As in the CAPM,

this assumption is equivalent to the firm having a utility function which is quadratic in profits or that the random input prices p_t have a multivariate Gaussian distribution. Thus, the firm's problem is, at every time period, to

$$\underset{q_t, s_t, y_t}{\text{maximize}} \; U[E(\pi_t'y_t - p_t'q_t - r_t's_t \mid I_t), V(\pi_t'y_t - p_t'q_t - r_t's_t \mid I_t)] \tag{2.1}$$
$$= U[\pi_t'y_t - r_t's_t - E(p_t'q_t \mid I_t), V(p_t'q_t \mid I_t)]$$
$$\text{subject to } f(y_t, s_t, Q_t) = 0, \quad \iota'q_t = Q_t$$
$$q_t, s_t, y_t \geq 0.$$

This optimization problem is equivalent to the two-stage process whereby first an optimal portfolio of suppliers is chosen to yield a given Q_t. Then, in the second stage, the proper balance is struck among outputs (y_t), non-risky inputs (s_t), and the total amount of the risky input (Q_t). The portfolio of q_t for a given Q_t and F (described below) is the solution to

$$\underset{q_t}{\text{maximize}} \; U(F - E(p_t'q_t \mid I_t), V(p_t'q_t \mid I_t)) \tag{2.2}$$
$$\text{subject to } \iota'q_t = Q_t, \quad q_t \geq 0.$$

Substituting this vector of optimal supplier quantities back into the objective function yields the optimal value function $U^*(F, Q_t \mid I_t)$ where F is net income from non-risky inputs and outputs. Thus U^* defines the highest level of utility obtainable for a given F and Q_t; the optimal q_t^* (which is a function of F and Q_t) has been substituted in for q_t. The second stage optimization problem uses this optimal value function to determine the utility maximizing total quantity of the risky input (Q_t), non-risky inputs (s_t) and outputs (y_t) as follows:

$$\underset{y_t, s_t, Q_t}{\text{maximize}} \; U^*(\pi_t'y_t - r_t's_t, Q_t \mid I_t) \tag{2.3}$$
$$\text{subject to } f(y_t, s_t, Q_t) = 0, \quad Q_t, s_t, y_t \geq 0.$$

Clearly, (2.3), with U^* defined by (2.2), is equivalent to solving (2.1).

Because we are only interested in the choice of the portfolio of suppliers of the risky

input, we will focus on (2.2). Ignoring the possible negativity of any elements of q_t, the Lagrangian for (2.2) is:

$$L = U(F - \mu_t'q_t, q_t'\Sigma_t q_t) + \eta(Q_t - \iota'q_t) \tag{2.4}$$

where η is the Lagrange multiplier on the constraint that the sum of purchases from various suppliers equals Q_t. The first-order conditions from (2.4) are:

$$\frac{\partial L}{\partial q_t} = -U_1\mu_t' + 2U_2 q_t'\Sigma_t - \eta\iota' = 0, \tag{2.5}$$

where U_i is the derivative of U with respect to its ith argument. Equation (2.5) can be solved for the scalar η using the constraint $\iota'q_t = Q_t$:

$$\eta = \frac{2U_2 Q_t - U_1\iota'\Sigma_t^{-1}\mu_t}{(\iota'\Sigma_t^{-1}\iota)}. \tag{2.6}$$

Substituting (2.6) back into (2.5) and re-arranging gives the following expression for the optimal vector of "risky" input shares:

$$w_t = \left[\frac{\lambda_t Q_t + \iota'\Sigma_t^{-1}\mu_t}{(\iota'\Sigma_t^{-1}\iota)}(\Sigma_t^{-1}\iota) - \Sigma_t^{-1}\mu_t\right]\frac{1}{\lambda_t Q_t} \tag{2.7}$$

where $\lambda_t = -\dfrac{2U_2}{U_1}$. Except for the 2, λ_t is the producer's marginal rate of substitution between expected costs and risk. It is, of course, a function of I_t, Q_t and F. However, the Q_t and F, and thus λ_t, are the result of solving (2.3). Rather than solve (2.3) explicitly, we make an assumption about λ_t. Several specifications for λ_t are possible. The first is simply $\lambda_t = \lambda$ for all t. Another, which is the specification we adopt, is that $\lambda_t = \dfrac{\lambda}{Q_t}$; i.e., $\lambda_t Q_t$ is a constant. This specification for λ_t has the attractive feature that it makes the optimal input share (2.7) invariant to Q_t. Thus as there is a secular rise in the level of productive activity in the firm, supplier shares remain constant. This expression for λ_t simplifies (2.7) to

$$\mathbf{w}_t = \left[\frac{\lambda + \iota'\Sigma_t^{-1}\mu_t}{(\iota'\Sigma_t^{-1}\iota)}(\Sigma_t^{-1}\iota) - \Sigma_t^{-1}\mu_t \right]\frac{1}{\lambda}. \tag{2.8}$$

For notational ease in what follows we write (2.7) as:

$$\mathbf{w}_t = S_t(\mu_t, \Sigma_t, \lambda) \tag{2.9}$$

Therefore given values for μ_t, Σ_t, and knowing its value of λ, the firm can compute the optimal period t input supplier mix from equation (2.8).

We assume that the firm knows or behaves as if it knows the parameters of the stochastic process determining the time path of the vector of input prices so that it can compute μ_t and Σ_t for all t. Unfortunately, in order for us to implement this model and test its empirical validity we must estimate the parameters of this stochastic process. Therefore, we now turn to the econometrics of the risk diversification model of input demand.

3. The Econometrics of the Risk Diversification Model of Input Demand

In this section we present our methodology for implementing the risk diversification model of input demand. There are two independent sources of uncertainty in this model. The first, what we call estimation error, arises from the estimation of μ_t and Σ_t, the conditional mean and covariance matrix of the vector-valued price process. The second, what we refer to as optimization error, is included to account for any unobservable time specific random shocks which may cause the first-order conditions (2.5) not to hold exactly each period.

This optimization error has the implication that we require the first-order conditions to hold only in expectation. Operationally, this means that (2.9) becomes:

$$\mathbf{w}_t = S_t(\mu_t, \Sigma_t, \lambda) + \epsilon_t \tag{3.1}$$

where $\epsilon_t \in R^n$ is $N(0,\Omega)$. The restriction that $\iota'\mathbf{w}_t = 1$ implies that $\iota'\epsilon_t = 0$ and $\iota'\Omega\iota = 0$. We can also interpret this optimization error as a way to take into account uncertainty in

input deliveries. In other words, on the average the firm correctly contracts from the various suppliers and on the average it actually receives these contracted amounts, so that $E(w_t) = S_t(\mu_t, \Sigma_t, \lambda_t)$. Hence, despite requiring the total input supply to be deterministic, the model does allow random variation in deliveries across suppliers from their utility maximizing levels.

We now discuss our methodology for modeling the estimation error in μ_t and Σ_t. There are two basic approaches available for estimating these parameters. The first approach is to model the price vector process using standard vector ARIMA techniques, allowing the data complete freedom to specify the way in which the past prices from other suppliers effect the present price of each supplier. However, the relatively high dimension of the price vector p_t makes this type of model very difficult to fit for data sets of the usual lengths found in applied econometrics. This is so because each additional autoregressive or moving average matrix added to the model requires the estimation of n^2 more parameters. For example, in our empirical application $n = 5$; thus in order to fit even a vector autoregressive process of order 1, there are 25 parameters to be estimated. Granted, some of these parameters may be zero, but determining which of them are zero is extremely difficult for the usual sample sizes found in economic time series. In addition, for this size n, the data set sizes required for the small sample reliability of asymptotic tests of the null hypothesis of multivariate white noise innovations from a vector ARIMA model are quite large (see Hosking (1980)). This further compounds the problem of reliably fitting a parsimonious vector ARMA model for this size n, because one of the usual requirements of model adequacy is the inability to reject the null hypothesis of white noise errors.

Our approach to modeling the price process still takes into account the interaction between the past prices other suppliers charge in determining each suppliers present price.

However, we impose structure on the model provided by economic theory to limit the number of parameters estimated. We postulate an error-correction model for each of the price processes. The basic idea behind these models is that there exists some stochastic equilibrium relationship between the elements of the vector-valued stochastic process in the sense that some linear combination of its elements has a time invariant mean and autocovariance function and the stochastic process has a long-run tendency towards making this linear combination equal its mean. In brief, this linear combination is a stationary univariate stochastic process. The tendency to return to "steady state" is accomplished by the model specifying that a proportion of the deviation of this linear combination of the prices from its steady state value in one period be corrected in .the next period. In our context this implies that the change in this period's price from supplier i is a function of the magnitude of this stationary linear combination of all of last period's prices. For instance, if this linear combination were a simple weighted average, a supplier's current price would be reduced if this average price last period exceeded its long-run mean. This class of models has been advocated and popularized in time series econometrics by Davidson, Hendry, Srba, and Yeo (1978). Various papers have also attempted to provide a more rigorous economic justification for the error-correction model by showing that it arises as the solution of a dynamic optimization problem (see Currie (1981) and Salmon (1982)).

Recent work by Engle and Granger (1987) has shown a close relationship between error correction models and the concept of co-integration. The Granger Representation Theorem of Engle and Granger (1987) states that any co-integrated process has an error correction representation. Although Engle and Granger (1987) provides an excellent discussion of co-integrated processes and error correction models, for continuity we restate their definitions here.

Definition. A series with no deterministic component which has a stationary invertible ARMA representation after differencing d times, is said to be integrated of order d, denoted $x_{it} \sim I(d)$.

Definition. The components of the vector x_t are said to be co-integrated of order d, b, denoted $x_t \sim CI(d,b)$, if (i) all components of x_t are I(d); (ii) there exists a vector $\alpha \neq 0$ so that $z_t = \alpha'x_t \sim I(d-b)$, $b > 0$. The vector α is called the co-integrating vector.

For our purposes we concentrate on the case that $d = b = 1$, so that the first difference of each of the series is stationary and some linear combination of the levels of all of the series is also stationary.

Definition. A vector time series x_t has an error correction representation if it can be expressed as:

$$A(B)(1 - B)x_t = -\gamma z_{t-1} + u_t \qquad (3.2)$$

where u_t is a stationary multivariate disturbance, B is the backshift operator, $A(0) = I$ (an n×n identity matrix), $A(1)$ has all elements finite, $z_t = \alpha'x_t$ and $\gamma \neq 0$.

The notation A(i) denotes the ith (n×n) matrix of coeffcients of the matrix polynomial in powers of the backshift operator B ($B^k x_t = x_{t-k}$) representing the vector autoregressive component of the process Δx_t. The process x_t still has a vector ARMA representation in levels as $\phi(B)x_t \equiv \psi(B)v_t$, where $v_t \sim N(0,\Omega)$. However, as discussed in Engle and Granger (1987), if a process has an error-correction representation this implies nonlinear restrictions on the parameters of the matrix polynomials in the ARMA representation. Part 3 of the Granger Representation Theorem discusses this issue.

Ignoring presence of co-integration in the estimation of vector ARMA models has several consequences. As stated by Engle and Granger (1987) ignoring the presence of co-integration and estimating a vector ARMA model in first differences leads to a specification

error because lagged levels of the variables are incorrectly left out of the model. To see this note that (3.2) is the data generating process and a vector ARMA model in first differences is the same as (3.2) but the term $-\gamma z_{t-1}$ is omitted. This omission results in the standard left out variables specification error. If a vector ARMA model in the raw series is estimated, important cross equation restrictions in the parameters of the model will not be imposed. Although these constraints will be satisfied asymptotically, efficiency gains will be obtained if they are imposed. In addition, the presence of roots of the matrix polynomial $\det(A(z)) = 0$ near unity will make precise estimation of individual parameters of the vector ARMA model in levels extremely difficult. Consequently, correctly detecting the presence of co-integration in a vector ARIMA model allows the imposition of cross-equation restrictions which leads to greater efficiency in estimation. However, rather than incorrectly imposing this restriction on our vector ARIMA model we should first test for its existence.

To test for co-integration in our price vector process we follow the procedure described in Engle and Granger (1987). First we test for a unit root in the backshift operator polynomial of the autoregressive portion of the univariate ARIMA representation of each price series using a Dickey-Fuller (1979) test. Then using this same test we verify the first-differences of each of the series *do not* have a unit root in the autoregressive polynomial portion of their univariate ARMA representation. The desired outcome of these two tests is that the first test in the raw series verifies the non-stationarity of this series and the second test in first-differences confirms the stationarity of the differenced series. This gives us the first requirement of co-integration: each of the price series is I(1). The second half of the co-integration testing procedure requires us to find a stationary linear combination of the price series. For this we run two of the tests in Engle and Granger (1987) for co-integration to verify that the vector of input prices are co-integrated: one based on the Durbin-Watson

statistic from the regression of one price on all of the other prices and the other is a modified Dickey-Fuller test on the residuals from this regression. Conditional on our test statistics confirming that p_t is co-integrated of order 1, we fit an error correction model for each of the input price series.

The general error-correction model fit for each price series is:

$$\Delta p_{it} = c_i + \gamma_i z_{it-1} + \beta_i \Delta p_{it-1} + \xi_{it} \quad (i = 1,...,n) \quad (t = 1,...,T) \tag{3.3}$$

where $\Delta p_{it} = p_{it} - p_{it-1}$ and $z_{it} \equiv \alpha_i' p_t$ is the estimate of the stationary linear combination of all of the prices (the scalar z_t defined earlier) for the ith price equation. Following Engle and Granger (1987) and Stock (1987), z_{it} is the residual computed from the regression of p_{it} on all other prices and a constant. Hence z_{it} has a sample mean of zero. As shown in Stock (1987), because the parameters of the co-integrating regression converge to their true values at rate T, rather than the usual $\sqrt{T}$, the z_{it} may be effectively treated as the observed z_t in the estimation of the parameters of (3.3) and the computation of their $\sqrt{T}$-asymptotic distribution.

We assume that the ξ_{it} is the ith element of ξ_t, which is distributed as a $N(0,\Sigma)$ random vector. This distributional assumption for ξ_t and the model (3.2) for p_{it} (i=1,...,n) implies that the conditional variance of Σ_t equals a constant Σ for all t.

Besides embodying the co-integration property of p_t, this model for each price series is consistent with the following economic logic. The constant term c_i takes into account the possibility that Δp_{it} may have a nonzero mean. By including this constant term we are effectively allowing p_{it} to have a deterministic trend. The second term in z_{it-1} is the error correction term which takes into account the fact that the amount z_{t-1} differs from its steady state value will effect this period t's price change for this supplier. We would expect γ_i to be negative to reflect the logic that if z_{it-1} is positive (recall the definition of z_{it}), this

period's Δp_{it} should be lower than its mean to reflect a correction towards steady state. The third term in Δp_{it-1} represents the impact of last period's price change on this period's price change. Model (3.2) represents the starting point for specifying the process followed by each price series. Additional terms in Δp_{js}, $(j = 1,...,n)$, $(s \leq t-1)$, are added to the model until the error process ξ_{it} resembles white noise, so that the final model estimated for each price series will be such that the null hypothesis that $\xi_{it} = E(p_{it} \mid I_t) - p_{it}$ is white noise cannot be rejected. As discussed earlier, in our case, the asymptotic tests for multivariate white noise require too many observations be of much use. Instead we concentrate our specification analysis on assuring that ξ_{it} is univariate white noise for each supply price process. Let Γ_i denote all of coefficients entering into (3.3) for Δp_{it}. Let $\mu_{it}(\Gamma_i, I_t)$ denote the conditional mean function for the ith price process. In this shorthand notation we can rewrite (3.3) as

$$p_{it} = \mu_{it}(\Gamma_i, I_t) + \xi_{it} \qquad (i = 1,...,n). \qquad (3.4)$$

If we stack all of the Γ_i into a single vector Γ then we can write (3.4) in vector notation as:

$$p_t = \mu_t(\Gamma, I_t) + \xi_t. \qquad (3.5)$$

Once we fit a univariate model to each price series such that the null hypothesis that each ξ_{it} series is white noise cannot be rejected, we can construct a consistent estimate of Σ as follows:

$$\hat{\Sigma} = \frac{1}{T} \sum_{t=0}^{T} \hat{\xi}_t \hat{\xi}_t{}',$$

where $\hat{\xi}_t = (\hat{\xi}_{1t}, \hat{\xi}_{2t}, ..., \hat{\xi}_{nt})'$ and $\hat{\xi}_{it}$ is the OLS estimate of ξ_{it}. This completes the first step of our two-step procedure to obtain consistent estimates of Σ and Γ. In summary, the equation by equation OLS estimates of the Γ_i in (3.4) yield a consistent estimate of Γ and because $\hat{\Sigma}$ is based on this estimate of Γ, it is also consistent.

The next step of this procedure conditions on these estimates of Γ and Σ and computes (from (3.1)) estimates of λ and Ω by maximum likelihood. We use maximum likelihood as opposed to OLS equation by equation for two reasons. The first is, OLS equation by equation would yield n different consistent estimates of λ, with no *a priori* reason to select one over the other. The second reason is that the maximum likelihood estimator is invariant to the equation dropped when estimating share equations, so that this procedure will yield the same estimate of λ taking into account the cross-equation restrictions on λ, regardless of which equation is dropped. Once this second stage is completed we have $\sqrt{T}$-consistent estimates of all of the parameters of our econometric model of the risk diversification approach to input demand. However, as we are also interested in testing this model we must also compute a consistent estimate of the asymptotic standard error of $\hat{\lambda}$, our second stage estimate of λ.

To compute a consistent estimate of the standard error of $\hat{\lambda}$ we need to know the asymptotic distribution of $\hat{\Gamma}$ and $\hat{\Sigma}$ our consistent estimates of Γ and Σ. The computation of the asymptotic distribution of Γ is fairly straightforward and can be obtained under more general distributional assumptions for ξ_t than multivariate normal. However, the asymptotic distribution of $\sqrt{T}\hat{\Sigma}$ explicitly depends on the distribution assumed for ξ_t and only takes a computationally tractible form in the case that ξ_t is multivariate normal or a member of the family of elliptical distributions. Consequently, in order compute a consistent estimate of the standard error of $\hat{\lambda}$ we must assume a distribution for ξ_t. We assume ξ_t is multivariate normal. This distributional assumption for the price process implies that our expected cost and variance of cost preferences for the firm result in no loss of generality, because the first two moments of the price process completely characterize its distribution.

Combining the model determining w_t in (3.1) with that determining p_t in (3.4) yields the following nonlinear maximum likelihood model:

$$\begin{bmatrix} p_t \\ w_t \end{bmatrix} = \begin{bmatrix} \mu_t(\Gamma, I_t) \\ S_t(\mu_t(\Gamma, I_t), \Sigma, \lambda) \end{bmatrix} + \begin{bmatrix} \xi_t \\ \epsilon_t \end{bmatrix}, \tag{3.7}$$

where $E(\xi_t \epsilon_t') = 0$ because the estimation error is independent of the optimization error by assumption.

The log-likelihood function is:

$$\ln L = -\frac{T(2n-1)}{2} \ln 2\pi - \frac{T}{2} \ln \det(\Sigma) - \tfrac{1}{2} \sum_{t=1}^{T} (p_t - \mu_t(\Gamma, I_t))' \Sigma^{-1} (p_t - \mu_t(\Gamma, I_t))$$
$$- \frac{T}{2} \ln \det(\Omega) - \tfrac{1}{2} \sum_{t=1}^{T} (w_t - S_t(\mu_t(\Gamma, I_t), \Sigma, \lambda)' \Omega^{-1} (w_t - S_t(\mu_t(\Gamma, I_t, \Sigma, \lambda)).$$

Given the two-step $\sqrt{T}$—consistent estimates of Γ, Σ, Ω, and λ described earlier, by the logic of Theorem 6.3.1 of Lehmann (1983) asymptotically efficient estimates of these parameters can be obtained by one iteration of a method of scoring type algorithm. Alternatively, starting from these consistent estimates and running this iterative procedure to convergence also yields asymptotically efficient estimates of these parameters.

We use the procedure suggested by Berndt, Hall, Hall, and Hausman (1974), hereafter BHHH, to compute the iterative maximum likelihood estimates of these parameters. To simplify the computational complexity of the problem we estimate Σ and Ω in terms of the parameters of the Cholesky decomposition of their inverses. Recall that Σ^{-1} can be written as LDL', where L is a lower triangular matrix with 1's along the diagonal and D is a diagonal matrix. The determinant of Σ^{-1} is the product of the diagonal elements of D. This decomposition simplifies the terms in the likelihood containing the determinant of Ω and Σ to a product of four diagonal elements in the former case and the product of five diagonal elements in the latter case. By the invariance property of maximum likelihood estimation, the

maximum likelihood estimates of Ω and Σ are equal to the inverse of the maximum likelihood estimates of the Cholesky decomposition of the parameters of their respective inverse matrices. Consistent estimates of the standard errors can be obtained from the sample average of the matrix of outer products of the gradients of the log-likelihood function evaluated at the maximum likelihood estimate of the parameter vector as described in BHHH. This estimation procedure uses as starting values the two-step consistent estimates and then it iterated to convergence.

Given this framework for specifying and estimating our model of input choice under price uncertainty we are now ready to apply it to the Japanese steam coal import market. Before proceeding to the application we first describe the history and operation of this market.

4. The Japanese Steam Coal Import Market

Almost immediately after the 1973-74 Arab Oil Embargo and subsequent substantial increase in the world price of oil, the Japanese embarked on a plan, co-ordinated between business and government, for stable domestic energy supply (Wu, 1977). Major among the methods Japan used to achieve this goal was to diversify both the suppliers and sources of energy. Previous to this event Japan had a oil based energy economy and it obtained most of this oil from the Middle East and United States. Subsequent to this embargo Japan expanded its sources of oil to China and the Soviet Union and began to consider coal as a major source of energy.

At the time Japan was importing coal primarily from the United States for use as coking coal in the production of steel. By the beginning of 1977, the Soviet Union, China, South Africa, and Australia had become consistent participants in this market, but the United

States was still the major producer of Japanese coal imports. By this time, Japan was also importing steam coal to be burned in coal-fired electricity generation facilities. Also, in rapid response to the oil price rises, Japan quickly converted most of its cement manufacturing plants from oil-fired to steam coal-fired (Tukenmez and Tuck, 1984). During the next five years the United States' share of the steam coal market steadily declined and the shares of South Africa and Australia increased considerably. The average volume of monthy steam coal imports (as classified by the Japan Tariff Association) rose from approximately 300,000 metric tons per month in early 1977, to 1.5 million metric tons per month in early 1984, and eventually to close to 2.7 million metric tons per month in mid-1987.

This steam coal is imported through negotiations with Japanese trading companies in conjuction with MITI for delivery to the steam coal using facilities. Prices for coal are negotiated in terms of the currency of the country of origin of the coal, although sometimes in dollars. Hence the price risk borne by Japanese consumers is primarily due to foreign exchange rate risk. Coal is purchased using three mechanisms: joint venture between buyer and supplier, long-term contract, and short-term supply agreement. While most of this coal is negotiated for delivery through long-term contracts, the Japanese trading companies often re-negotiate these contracts when current market conditions favor their doing so. For example, when there is a downturn in the world coal market many of these contracts are re-negotiated. This potential for re-negotiation of long-term contracts based on current market conditions is another source of price uncertainty.

There is an abundance of anecdotal evidence for the validity of the risk diversification model of input choice for the Japanese steam coal import market. In various editions of the MITI Handbook published by the Japan Trade and Industry Publicity, Inc., two major policy goals for MITI in the area of energy and natural resources are: (1) a stable supply of energy

resources and (2) stable prices of energy resources. One of the stated goals of the Coal Mining Department of MITI is "to smooth the importation of coal," (MITI Handbook 1979/80, p. 82). Japan's desire for a stable, secure energy supply is well-documented in Wu (1977), a recent study of Japan's response to the Arab Oil Embargo of 1973/74. In addition, a recent U.S. Department of Energy study of coal trade in the Asian market states, "...in seeking diversification and security Japan seems willing to pay a premium to access stable coal supplies from the more expensive exporters, such as the United States...," (Tukenmez and Tuck (1984)). This casual evidence coupled with the three puzzles concerning the time series properties of the prices and quantities of imports of steam coal to Japan stated in the introduction makes for a challenging application of our risk diversification model of input choice that is also of substantial policy interest.

5. Application to Japanese Steam Coal Import Market

Time series of prices and quantities of steam coal[3] imported into Japan from China, Soviet Union, United States, South Africa, and Australia are available on a monthly basis from the Japanese Foreign Trade Statistics compiled by the Japan Tariff Association. All prices are in units of thousands of Yen per metric ton. The quantity units are metric tons. The input choice problem is invariant to the absolute price level. The normalization of prices will only effect the magnitude of λ. To make shares and prices of approximately the

---

[3]Classified by the Japan Tariff Association as "high and low ash coal other than coking coal." Although strictly speaking, steam coal from two different countries is two different commodities, steam coal is primarily, if not exclusively, valued for its heat content. Consequently, only coal with the highest heat content is exported. Although the heat content of each shipment of coal to Japan during the sample was not available, the heat content of coal of a representative sample of coal contracts from each of the supplier countries considered in this paper was available (TEX Report, Tokyo, 1986). For this representative sample, the mean coal contract heat content per ton of coal was not significantly different across the supplier countries considered here. This provides support for our treatment of steam coal from various countries as a homogeneous product.

same magnitude, prices were normalized so that the sum of the sample means of all prices of coal is equal to one. A convenient normalization for the quantities was in units of millions of tons. The sample period March of 1983 to May of 1987 was selected because the structure of the Japanese steam coal import market seems stable over this period. Confirmation of this point is that despite a growing total quantity of steam coal imported, the shares of this market served by each of the suppliers show no statistically significant serial correlation or trend over this period. This empirical fact provides further support for our selection of a form for λ_t which makes the optimal supplier shares independent of Q_t.

The first step of the estimation procedure is to test for co-integration between the five price processes over the sample. To confirm that each of the univariate price processes are integrated of order one, Dickey-Fuller (1979) tests for unit roots are performed on the levels and first differences of each series. The models run for each test are:

$$\Delta p_{it} = \alpha + \beta_1 p_{it} + \beta_2 \Delta p_{it-1} + e_{it} \tag{5.1}$$

for the test for a unit root in the levels, and

$$\Delta^2 p_{it} = \alpha + \beta_1 \Delta p_{it} + \beta_2 \Delta^2 p_{it-1} + d_{it} \tag{5.2}$$

for the test for a unit root in the first differences. In both cases the null hypothesis is that $\beta_1 = 0$, or more precisely, the backshift operator polynomial of the AR portion of the ARIMA representation of x_t (x_t represents either the raw or first-differenced price series) has the following factorization $\phi(B) = (1 - B)\phi^*(B)$ where all of the roots of $\phi^*(z) = 0$ are greater than one in modulus. The results of these tests are given in Table 5.1. For all of the tests in terms of the levels of the price series, there is little evidence against the null hypothesis of a unit root, indicating that nonstationarity of the price series in levels cannot be rejected. In constrast, the null hypothesis of a unit root in the first differenced series is decisively rejected for all of the series at 0.01 level of significance, providing strong evidence

for the stationarity of the first-differenced series. The critical value for the test is from Table 8.5.2 of Fuller (1976).

Table 5.1: Dickey-Fuller Tests for Unit Roots		
Critical value for 0.01 size test = -3.58		
Country	Levels	First-Differences
China	-0.8871	-6.0236
Soviet Union	-0.5911	-6.0028
United States	-0.8124	-7.2048
South Africa	-0.0488	-5.4708
Australia	0.4913	-5.6874

The results of this battery of tests is in line with the first requirement for the price processes to be cointegrated. The results of this table suggest that each of the univariate price processes is integrated of order 1.

The second requirement is that some linear combination of the prices is stationary. To test this hypothesis we perform two tests suggested by Engle and Granger (1987). The first test is performed for its ease of implementation as an intuitive informal test of the hypothesis and the second is performed because these authors recommend it as the best test available for co-integration. The first test is based on the Durbin-Watson statistic from the regression of p_{it} on a constant and the p_{jt} ($j \neq i$). As described earlier, this regression is called the co-integrating regression for the ith supplier and its residual vector is z_{it} used in (3.2). The inituition behind this test is that if the series are co-integrating, then the errors from this regression should be stationary. Therefore, the Durbin-Watson statistic should not detect a departure from this condition. If the Durbin-Watson statistic is close to zero, indicating non-stationarity of the errors, the null hypothesis of non-co-integration cannot be rejected. This test rejects the null hypothesis of non-co-integration in favor of co-integration if the Durbin-Watson statistic is too large. The second test is the augmented Dickey-Fuller

test (ADF). It is implemented via a Dickey-Fuller test in the form (5.1) given above on the residuals from the co-integrating regression for p_{it}. The null hypothesis of a unit root in the residual process corresponds to non-co-integration and the alternative of stationarity of the residual process corresponds to co-integration of the price processes. Table 5.2 contains the results of these test statistics and critical values. As can be seen from the table, the ADF tests on the residuals from the co-integrating regressions for China, the United States, and South Africa, imply that the null hypothesis that the series are non-co-integrating is rejected in favor of the alternative that they are co-integrating at the 0.01 level. These tests on the residuals from the regressions for Australia and the Soviet Union, find some evidence for co-integration, but the null hypothesis of non-co-integration cannot be rejected at a 0.05 level for either of these regressions.

Table 5.2: Regression Tests for Co-integration		
Country	Durbin-Watson	Augmented Dickey-Fuller
China	1.4606	-6.0708
Soviet Union	2.0097	-4.0123
United States	2.3619	-5.6352
South Africa	2.0355	-5.7883
Australia	1.4199	-3.4231
0.01 Critical Value	0.511	-4.80
0.05 Critical Value	0.386	-4.15

The critical values for the Durbin-Watson based tests were obtained from Table II of Engle and Granger (1987). These critical values were computed for the case that n=2. As discussed in Engle and Yoo (1987), the applicability of these critical values to the case considered here of n=5 is questionable, so these test statistics are presented primarily for their intuitive appeal. However, the critical values for the ADF statistics are those for the case n=5 from Table 3 of Engle and Yoo (1987). Nevertheless, the results of these two sets of tests provide strong evidence for the hypothesis that the five price series are co-integrating.

These results support the use of (3.3) to model each price series.

For each first differenced price series, the model given in (3.3) with a constant term, z_{it-1} and Δp_{it-1} was sufficient to adequately represent the behavior of each of the price processes over sample and still not reject white noise errors. Table 5.3 contains the results of these regressions. As expected, the signs of all of the parameters associated with z_{it-1} are negative. Because of the presence of z_{it-1}, the usual univariate Box-Pierce statistic for autocorrelation is not valid; instead the auxiliary regression form of Durbin's (1970) LM test for AR(1) disturbances was computed. For all of the models there is very little evidence for this alternative against the null hypothesis of univariate white noise errors. Generalizations of this test against general fourth order AR and MA processes were also performed but the null hypothesis of white noise errors could not be rejected for these cases as well.

Table 5.3: First Round Estimates of Price Processes				
Country	c_i	γ_i	β_i	t-test for AR(1) errors
China	-0.0022516	-0.71658	0.34724	-0.520287
S.U.	-0.0026191	-0.74900	-0.14908	0.137535
U.S.	-0.0035828	-1.16960	-0.04572	-0.821971
S.A.	-0.0028731	-0.68449	0.104030	-0.638958
Australia	-0.0032582	-0.12764	-0.34435	0.617483

Conditional on these first round estimates of the parameters of the price process, we then estimate λ and Ω by maximum likelihood. Having obtained these $\sqrt{T}$-consistent estimates of λ and Ω, we then compute fully efficient maximum likelihood estimates which impose the cross-equation restrictions implied by our risk diversification model and use these first round estimates as starting values. Table 5.4 contains the converged maximum likelihood estimate of Σ, the conditional covariance matrix of the price process.

Table 5.4: Maximum Likelihood Estimate of $\Sigma \times 10^4$					
Country	China	Soviet Union	United States	South Africa	Australia
China	3.0	1.2	0.6	1.2	-0.08
Soviet Union	1.2	10.6	-3.2	2.2	0.4
United States	0.6	-3.2	2.6	-0.3	-0.2
South Africa	1.2	2.2	-0.3	1.3	0.3
Australia	-0.08	0.4	-0.2	0.3	1.0

All maximum likelihood parameter estimates were within two standard errors (using the consistent standard errors estimates computed from the converged maximum likelihood parameter estimates as described at the end of section 3) of the first round set of consistent estimates of Γ, Σ, λ, and Ω, lending some credence to our maximum likelihood estimation procedure which jointly estimates the parameters of the price vector process, the parameters λ and Ω and imposes the cross equation restrictions implied by our structural model. These maximum likelihood parameter estimates and the consistently estimated standard errors allow a test of the suitability of the risk diversification approach to input demand. Table 5.5 contains the ML estimate of λ and its standard error.[4]

Table 5.5: ML estimate of λ	
$\lambda \times 10^{-3}$	1.123
Standard Error	0.827
t-statistic	7.365

Based on the asymptotic one-sided t-test H: $\lambda = 0$ versus K: $\lambda \geq 0$, we can reject the null hypothesis that λ is zero in favor of the alternative that it is positive. This hypothesis test provides support for the risk diversification model of input demand. In other words, based on this hypothesis test, Japan seems to attach a positive weight to the conditional variance of total input costs in choosing its optimal input supplier mix.

---

[4]Other maximum likelihood parameter estimates are not reported because of their agreement with, modulo two standard errors, the first round estimates in Table 5.3.

6. Implications of Risk Diversification Model of Input Demand

We now proceed to examine the implications of this estimated model of short-run input demand under price uncertainty. We are concerned with two general questions. The first is the size of the risk premium associated with steam coal imported to Japan. The second issue is how well does this model of input choice explain the three time series characteristics of supplier shares of steam coal imported to Japan discussed in section 1.

We first consider the question of the size of the risk premium on imported coal. To derive this magnitude consider the mean price versus standard error of price frontier plotted in Figure 1. Define $P_{p_t}(w) = w'p_t = \sum_{i=1}^{5} w_i p_{it}$ as the actual weighted average price of steam coal imports at time t, where $\sum_{i=1}^{5} w_i = 1$. Let $E(P_{p_t}(w))$ equal the expectation conditional on I_t of $P_{p_t}(w)$ and $\sigma^2(P_{p_t}(w))$ equal its variance conditional of I_t.[5] The mean-standard error frontier given in Figure 1 is comprised of the set of ($E(P_{p_t})$, $\sigma(P_{p_t})$) pairs such that $\sigma(P_{p_t}(w))$ is minimized over w subject to the constraints that $\iota'w = 1$ and $E(P_{p_t}(w)) = K$, where K is some positive constant. Once a value of λ is specified, the solution to (2.2) implies a point on the mean-standard error frontier and corresponding to the optimal input mix conditional on μ_t and Σ_t. This point is labelled $M(\hat{\lambda})$ in the diagram, with expected price $E(P_{mt})$ and standard error of price $\sigma(P_{mt})$. The slope of the mean-standard error frontier at the point $M(\hat{\lambda})$ is the rate at which Japan substitutes decreases in expected price for increases in the standard error of price at time t. Where the tangent line to the point $M(\hat{\lambda})$ intersects the expected price axis is represents the equilibrium expected price Japan would be willing to pay for riskless coal for a given $\hat{\lambda}$ at time t. We denote this price P_{st} because of the analogy to

---

[5]For the remainder of this section all expectations and variances are conditional on the information set at time t.

the zero-beta portfolio in the Capital Asset Pricing Model with no riskless asset.

To compute the portfolio of suppliers (weighted average price) which has no market risk (zero-beta) we solve for the minimum variance supplier weighted average price subject to the constraint that its covariance with the market price (P_{mt}) is zero. The Lagrangian for this optimization problem takes the form:

$$L = \tfrac{1}{2}w_{zt}'\Sigma w_{zt} + \eta(\tfrac{1}{2}w_{zt}'\Sigma w_{mt}) + \nu(1 - \iota'w_{zt}), \tag{6.1}$$

where w_{zt} is the independent variable, w_{mt} is the vector of optimal supplier shares from (2.6) ($P_{mt} = w_{mt}'p_t$) and η and ν are Lagrange multipliers associated with the constraints that the covariance of P_{zt} with P_{mt} is zero and that $\iota'w_{zt}$ is equal to one (ignoring the possible negativity of elements of w_{zt}). The solution to this optimization problem is:

$$w_{zt} = \frac{w_{mt} - w_{mt}'\Sigma w_{mt}(\Sigma^{-1}\iota)}{1 - (w_{mt}'\Sigma w_{mt})(\iota'\Sigma^{-1}\iota)}. \tag{6.2}$$

Because μ_t varies over time, for each t there is a different Figure 1 and corresponding P_{mt} and $P_{zt} = w_{zt}'p_t$. We define the risk premium at time t (RP_t) as

$$RP_t = \frac{E(P_{zt}) - E(P_{mt})}{E(P_{mt})}, \tag{6.3}$$

where P_{it}, i=z,m is the value of P_i at time t. Figure 2 contains a time series plot of RP_t based on the ML estimates of Γ, Σ, and λ. This risk premium ranges from 29 percent to 50 percent over the sample period, implying that Japan seems willing to pay from 29 to 50 percent above the current market price for a supply of coal having no price risk. This risk premium exhibits an increasing time trend. A risk premium that increases with time is consistent with the view that as Japan becomes more and more dependent of foreign sources of steam coal, as has been the case in recent years, the amount above the current weighted average market price Japan is willing to pay for riskless coal supply should increase.

We now turn to the issue of how well our model of input demand explains the time path of Japan's steam coal import shares. In order to address these issues we first describe one further implication of our risk diversification model of input choice. From equation (3.1) we know that the expected value of the observed vector of market shares is equal to the optimal vector of market shares based on μ_t, Σ and λ. Mathematically, this statement implies

$$E(w_t) = S(\mu_t, \Sigma, \lambda). \tag{6.4}$$

This condition implies the existence of an equilibrium market specific measure of risk for each supplier analoguous to the market specific measure of risk for each security in the Capital Asset Pricing Model (CAPM). For this reason we denote the market specific measure of risk for supplier i in period t by β_{it} and define it as

$$\beta_{it} = \frac{\text{Cov}(P_{mt}, p_{it})}{\text{Var}(P_{mt})} \tag{6.5}$$

where $P_{mt} = w_{mt}'p_t$ is the market price and p_{it} is supplier i's price. The covariance and variance in the expression for β_{it} are conditional on I_t, the information set at time t. Consequently, because the composition of P_{mt} will change each time period as μ_t changes, both the numerator and denominator of β_{it} will vary over time. Hence β_{it} will also change over time. Figure 3 contains the plot of the β_{it} for all suppliers over the sample period. Recall that by definition, P_{mt} has a β_t of one for all t.

Using logic similar to that used to derive the Security Market Line in the CAPM, we can derive an equilibrium relationship between the β_{it} and $E(p_{it} \mid I_t)$ as follows:

$$E(p_{it} \mid I_t) = E(P_{st} \mid I_t) + [E(P_{mt} \mid I_t) - E(P_{st} \mid I_t)]\beta_{it}, \tag{6.6}$$

where $E(P_{st} \mid I_t)$ is the conditional expectation of P_{st} and $E(P_{mt} \mid I_t)$ is the conditional expectation of P_{mt}.

We are now in a position to address the three puzzles presented in the introduction. The first puzzle is why the United States remains in the market despite its consistently high price. Figure 3 shows that the United States consistently has the lowest market specific measure of risk associated with it. In fact, in some periods beta for the United States is negative indicating that it is a good hedge against variations in the market price of coal. This characteristic of the price of United States coal explains the fact that its price path always lies above those of the other four countries and it has the second largest conditional variance (see Table 5.4), but it still services a sizeable share of this market. Furthermore, the negative elements in the United States' row and column in Σ explains the usefulness of the United States as a hedge supplier.

The second puzzle is why the Soviet Union is consistently the cheapest supplier but never captures much of the market. Figure 3 also shows that the Soviet Union consistently has the highest market specific measure of risk. In addition, from Table 5.4 the Soviet Union price has the highest conditional variance. These two risk measures illustrate why the Soviet Union has the smallest market share despite having the lowest price in most all periods. From equation (6.6) we can see that the high level of market specific risk associated with this supplier must be compensated for in terms of a low expected supply price in order for Japan to have a nonzero demand for this coal.

The last puzzle concerns why South Africa and Australia have similar prices but very different market shares. This can be answered by inspection of our estimate of Σ in Table 5.4. Australia has the smallest conditional variance in price and more importantly its price has virtually no conditional covariance with any of the other prices. Both of these points imply that its market share should be substantially larger than that of South Africa which has a higher conditional variance and has higher conditional covariances with the other

suppliers besides Australia. Finally, the similar time series behavior of the beta's associated with Australia and South Africa explain, in part, why the two price processes from these countries are very similar and the sample averages of the two price series are essentially the same.

7. Conclusions and Policy Implications

The risk diversification model of input demand seems to provide a useful framework for making economic sense of several puzzling anamomalies in the Japanese steam coal import market. Clearly, there are other models and factors which could explain the observed market shares; however, as mentioned earlier, the substantial anecdotal evidence for the applicability of the risk diversification model of input demand makes an examination of its validity of particular interest and relevance.

The policy implications of these results for suppliers of Japanese steam coal imports fall into two broad categories. The first, perhaps more naive view of these results, is that because the Japanese seem willing to pay a premium for stable prices, a country interested in supplying more of its coal to Japan should attempt to stabilize its price of coal in Yen to Japan. This view ignores the fact that much of the price uncertainty in coal to Japan is due to factors beyond the control of coal suppliers within that country. Supply interruptions, domestic price inflation in the country of origin, or price inflation in Japan all effect the price of coal in Yen in Japan. Consequently, perhaps a more sophisticated view of these results is that so long as each supplier's price process has some component of its variation which is independent of the variation in the prices of other suppliers, this supplier should have a nonzero market share so long as its prices are not too high above the prices of the rest of the suppliers.

Perhaps the most significant result to come out of our paper is the development of a rigorous but implementable methodology for representing input demand under price uncertainty and testing for risk diversification behavior in that framework. Future applications of this risk diversification model of input demand are plentiful. Any industry in which a large portion of variable costs is taken up by a single homogeneous factor of production represents a potential test of the risk diversification approach to input demand.

References

Batra, R.N. and A. Ullah, (1974), "Competitive Firm and the Theory of Input Demand under Price Uncertainty," *Journal of Political Economy,* 82, 537-548.

Berndt, E.K., B.H. Hall, R.E. Hall, and J.A. Hausman, (1974), "Estimation and Inference in Nonlinear Structural Models," *Annals of Economic and Social Measurement,* October, 653-665.

Blair, R.D., (1974), "Random Input Prices and the Theory of the Firm," *Economic Inquiry,* 12, 214-226.

Currie, D. (1981), "Some Long-Run Features of Dynamic Time-Series Models," *The Economic Journal,* 91, 704-715.

Davidson, J.E.H., D.F. Hendry, F. Srba, and S. Yeo (1978), "Econometric Modelling of the Aggregate Time-series Relationship Between Consumer's Expenditure and Income in the United Kingdom," *Economic Journal,* 88, 661-692.

Dickey, D.A. and W.A. Fuller (1979), "Distribution of the Estimators for Autoregressive Time Series with a Unit Root," *Journal of the American Statistical Association,* 74, 427-431.

Durbin, J.M. (1970), "Testing for Serial Correlation in Least Squares Regression when Some of the Regressors are Lagged Dependent Variables," *Econometrica,* 38, 410-421.

Engle, R.F. and C.W.J. Granger (1987) "Co-Integration and Error Correction: Representation, Estimation, and Testing," *Econometrica,* 55, 251-276.

Engle, R.F. and B.S. Yoo (1987) "Forecasting and Testing in Co-integrated Systems," *Journal of Econometrics,* 35, 143-160.

Fuller, W.A. (1976), *Introduction to Statistical Time Series,* John Wiley & Sons: New York.

Hosking, J.R.M. (1980), "The Multivariate Portmanteau Statistic," *Journal of the American Statistical Association,* 75, 602-608.

Lehmann, E.L. (1983) *Theory of Point Estimation,* John Wiley & Sons: New York.

Ministry of International Trade and Industry (1979), *The MITI Handbook,* Japan Trade and Industry Publicity, Inc.

Salmon, M. (1982), "Error Correction Mechanisms," *The Economic Journal,* 92, 615-629.

Sandmo, A. (1971), "Competitive Firm under Price Uncertainty," *American Economic Review,* March, 65-73.

Stock, J.H. (1987), "Asymptotic Properties of Least Squares Estimators of Cointegrating

Vectors," *Econometrica,* 55, 1035-1056.

TEX Report Ltd. (1986) *1986 Coal Manual,* Chiyoda-ku, Tokyo (1986).

Tuckenmez E. and N. Tuck (1984), "Coal-Exporting Countries: The Asian Market," U.S. Department of Energy Report, DOE/EIA-0462.

Wu, Y. (1977), *Japan's Search for Oil,* Hoover Institution Press: Stanford, California.
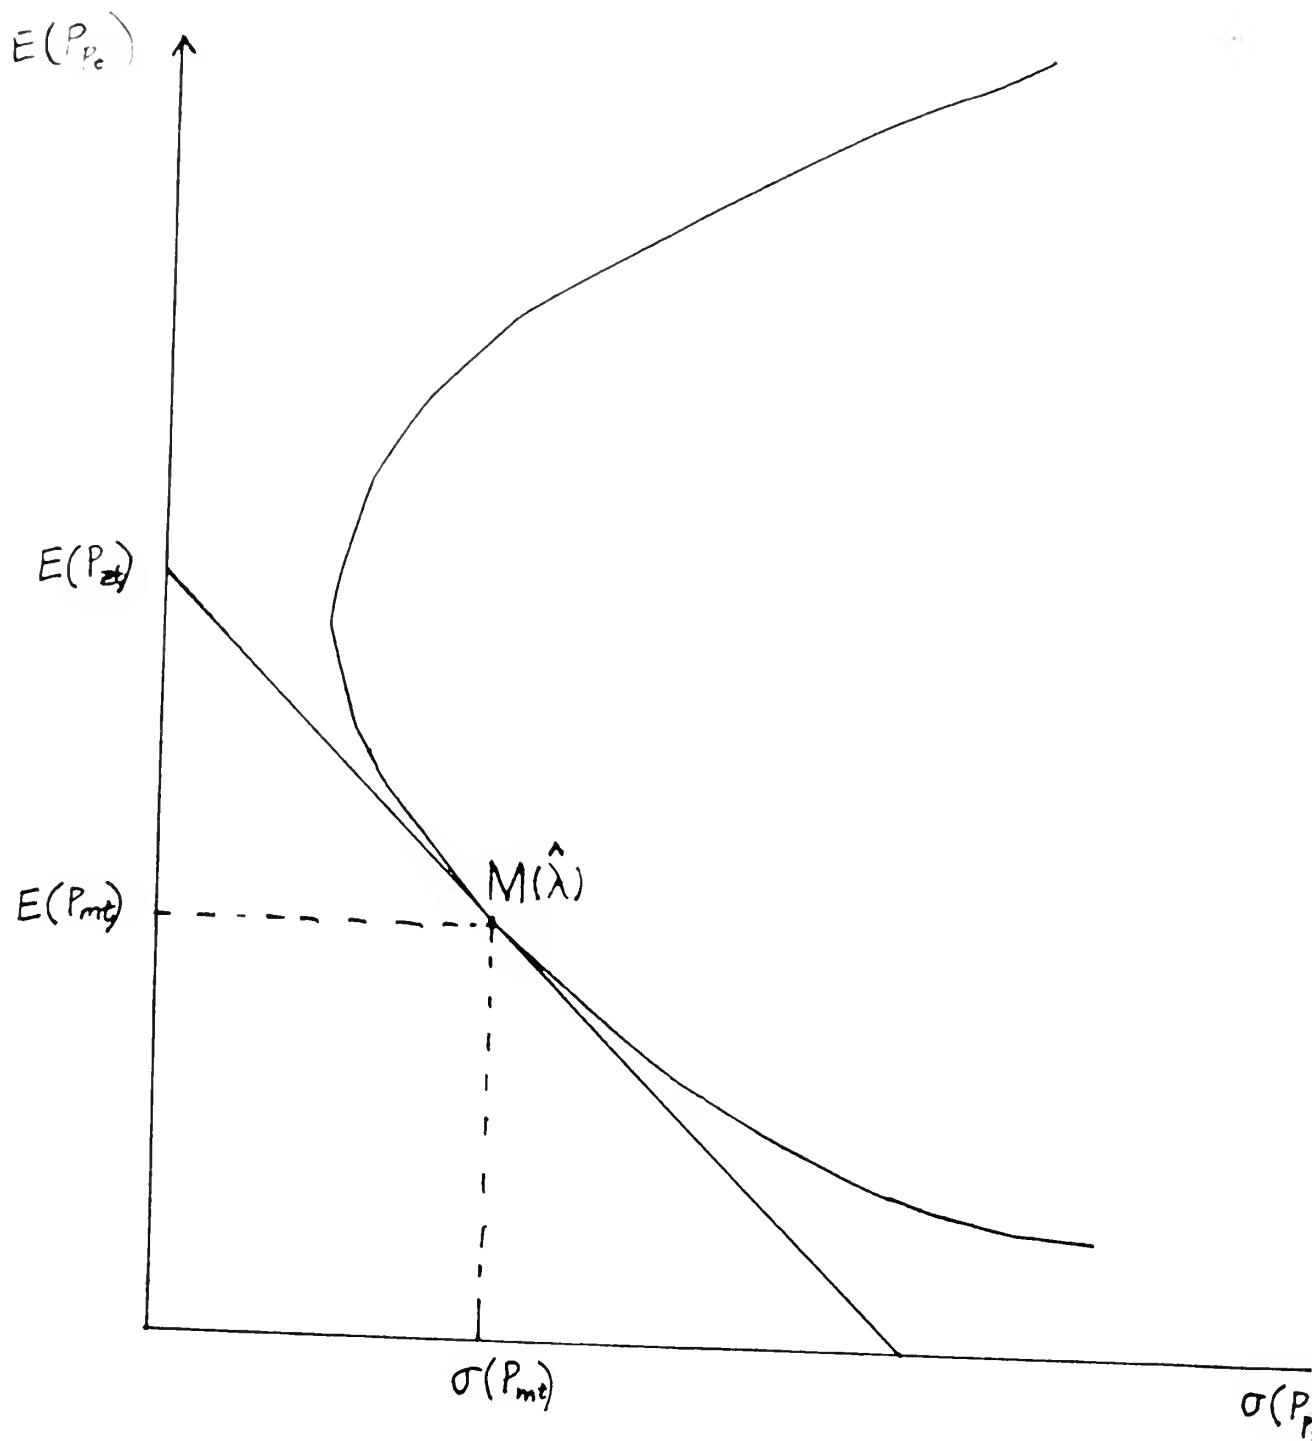
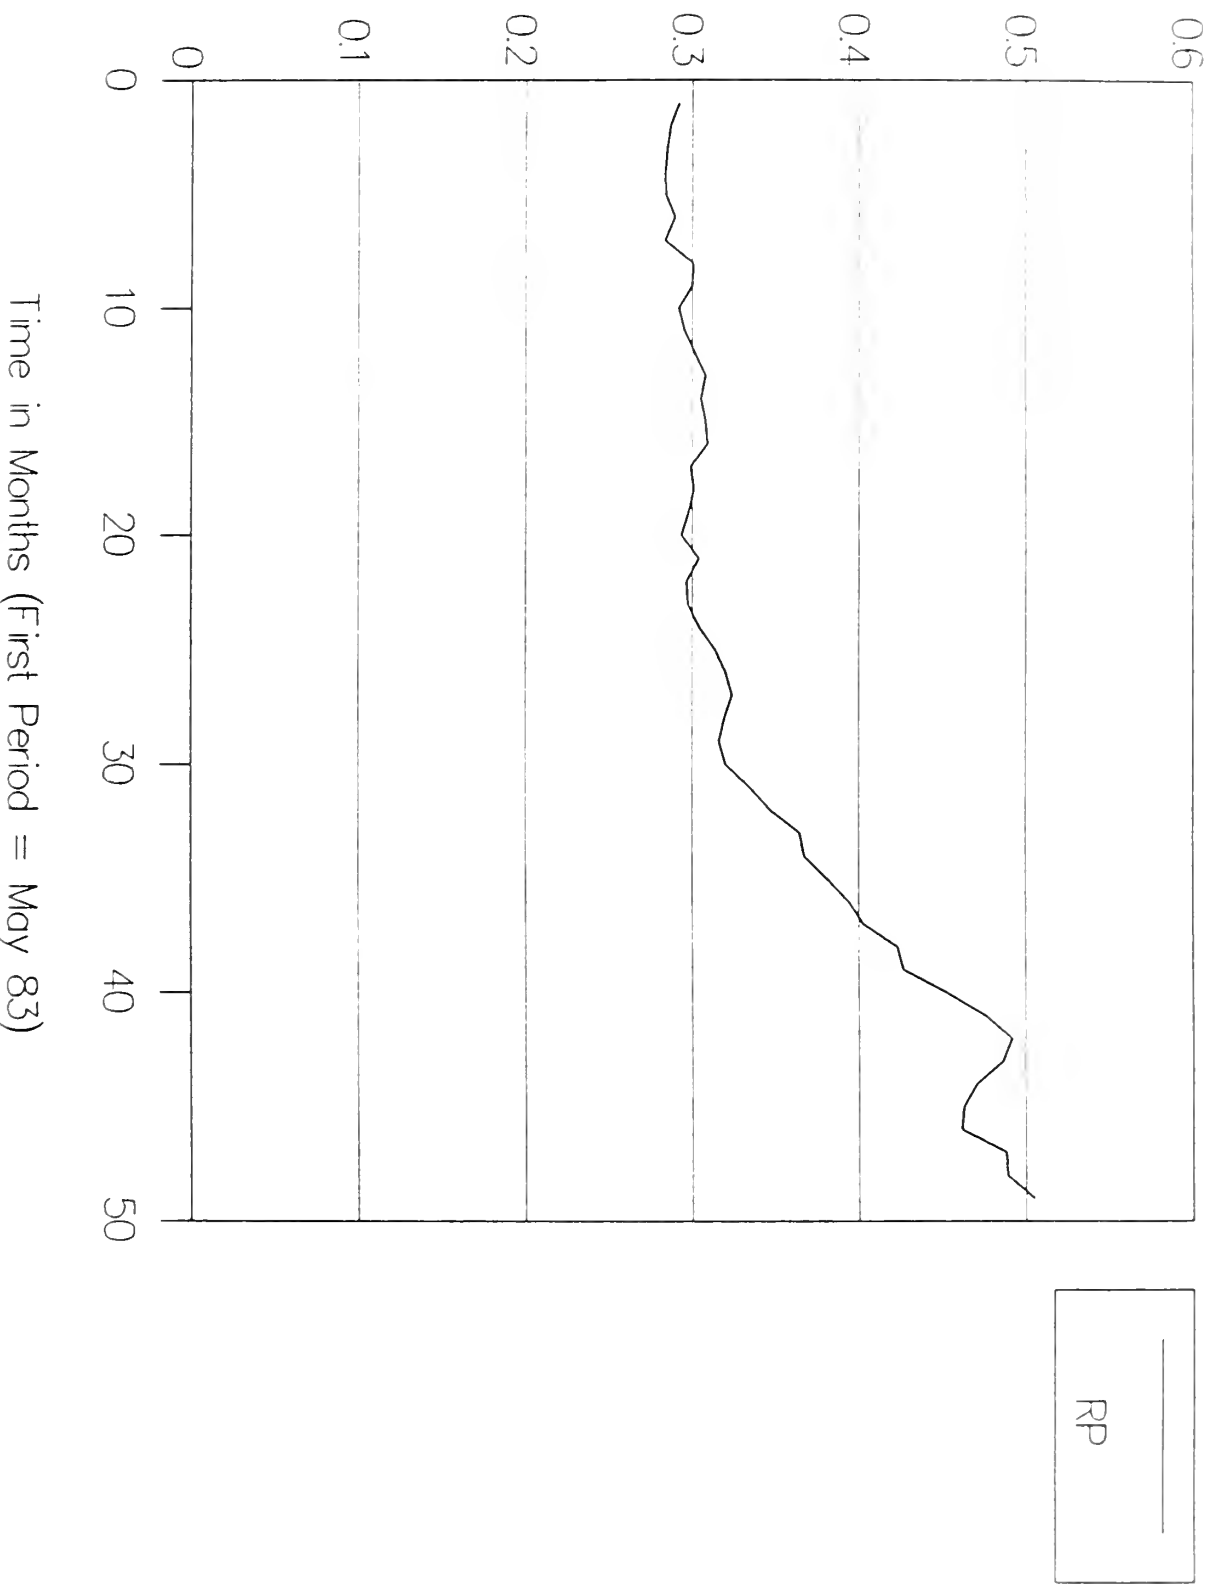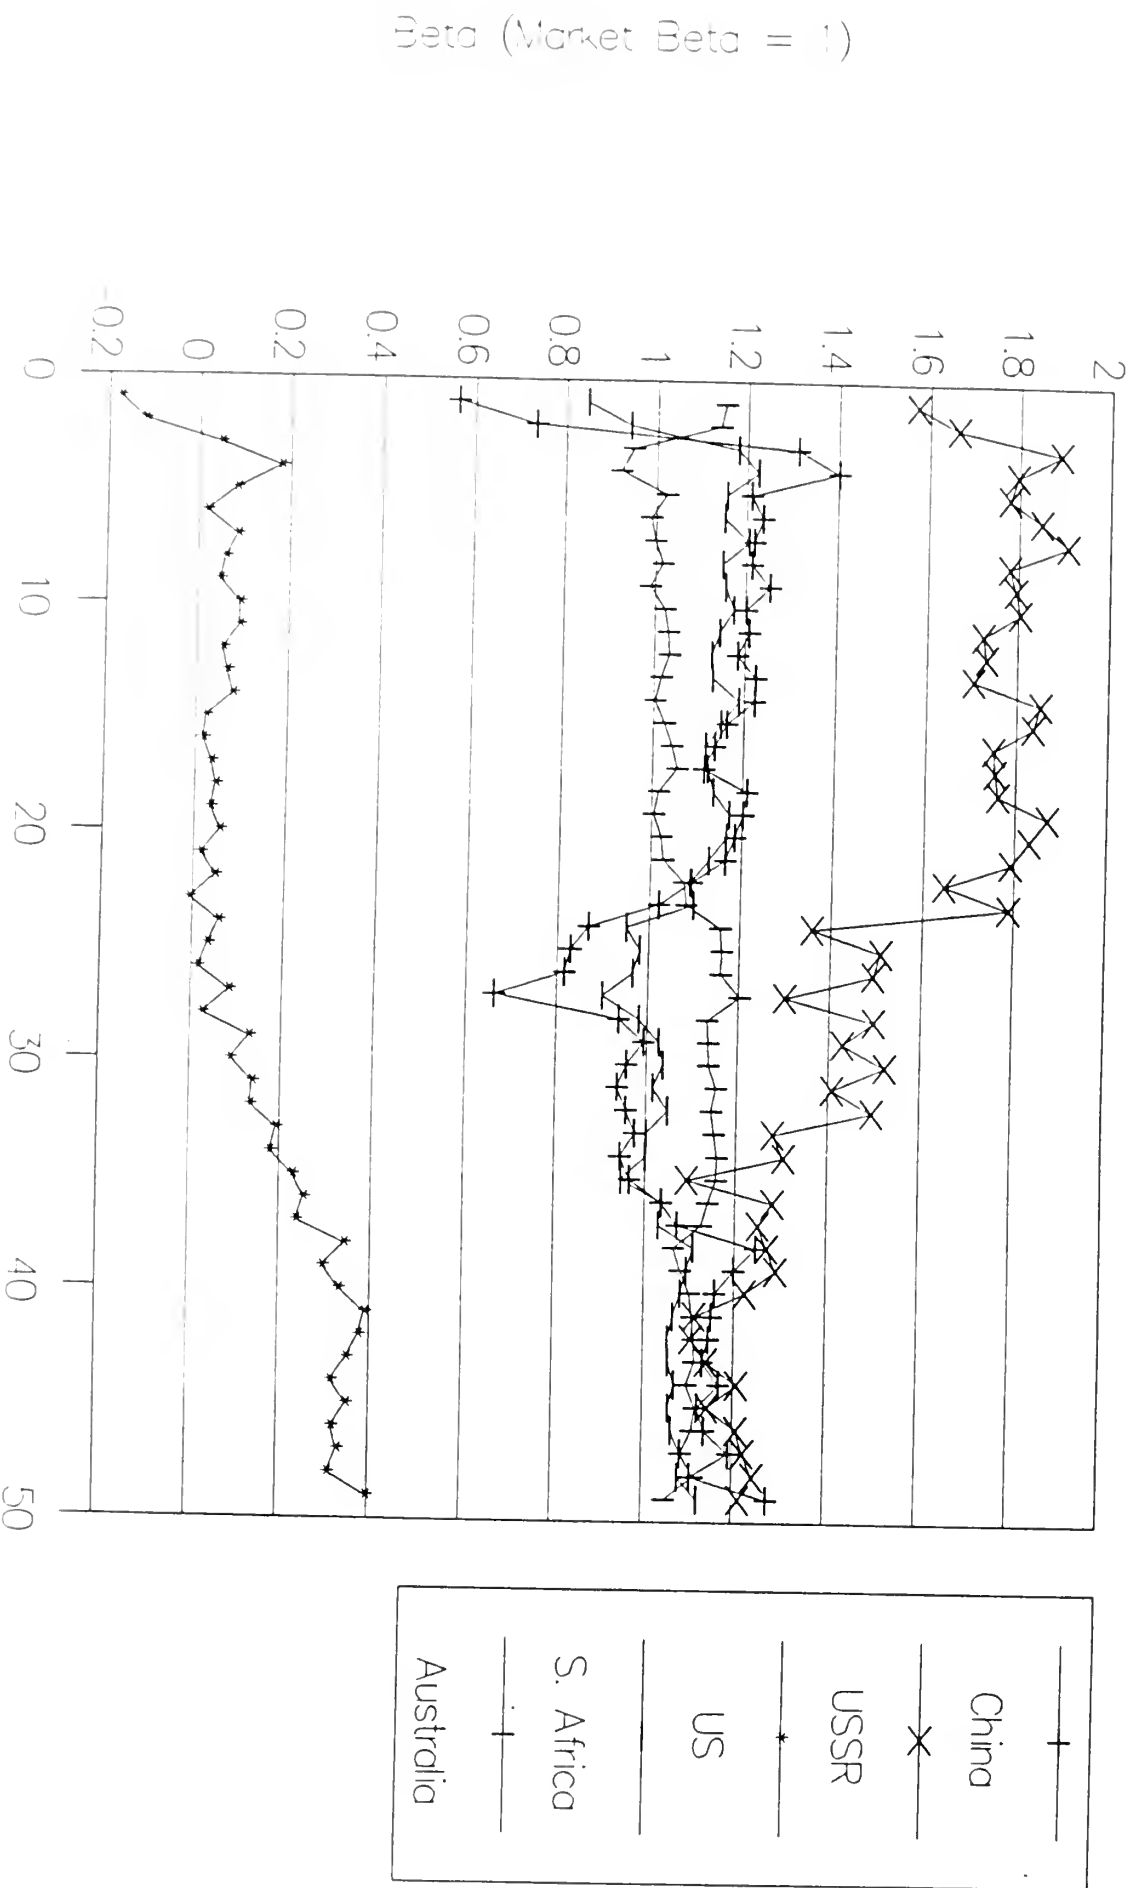FIGURE 1: EFFICIENT FRONTIER

Figure 2: Risk Premium for Sample

Figure 3: Betas for Sample Period

Beta (Market Beta = 1)

Time in Months (First Period = May 83)

China
USSR
US
S. Africa
Australia

Papers in the Political Economy of Institutions Series

No. 1   Susan I. Cohen. "Pareto Optimality and Bidding for Contracts."
         Working Paper # 1411

No. 2   Jan K. Brueckner and Kangoh Lee. "Spatially-Limited Altruism,
         Mixed Clubs, and Local Income Redistribution." Working Paper
         #1406

No. 3   George E. Monahan and Vijay K. Vemuri. "Monotonicity of
         Second-Best Optimal Contracts." Working Paper #1417

No. 4   Charles D. Kolstad, Gary V. Johnson, and Thomas S. Ulen. "Ex
         Post Liability for Harm vs. Ex Ante Safety Regulation:
         Substitutes or Complements?" Working Paper #1419

No. 5   Lanny Arvan and Hadi S. Esfahani. "A Model of Efficiency Wages
         as a Signal of Firm Value." Working Paper #1424

No. 6   Kalyan Chatterjee and Larry Samuelson. "Perfect Equilibria in
         Simultaneous-Offers Bargaining." Working Paper #1425

No. 7   Jan K. Brueckner and Kangoh Lee. "Economies of Scope and
         Multiproduct Clubs." Working Paper #1428

No. 8   Pablo T. Spiller. "Politicians, Interest Groups, and
         Regulators: A Multiple-Principals Agency Theory of Regulation
         (or "Let Them Be Bribed." Working Paper #1436

No. 9   Bhaskar Chakravorti. "Asymmetric Information, 'Interim'
         Equilibrium and Mechanism Design." Working Paper #1437

No. 10   Bhaskar Chakravorti. "Mechanisms with No Regret: Welfare
         Economics and Information Reconsidered." Working Paper #1438

No. 11   Bhaskar Chakravorti. "Communication Requirements and Strategic
         Mechanisms for Market Organization." Working Paper #1439

No. 12   Susan I. Cohen and Martin Loeb. "On the Optimality of
         Incentive Contracts in the Presence of Joint Costs." Working
         Paper #1445

No. 13   Susan I. Cohen and Martin Loeb. "The Demand for Cost
         Allocations: The Case of Incentive Contracts Versus Fixed-
         Price Contracts." Working Paper #1455

No. 14   Jan K. Brueckner and Kevin M. O'Brien. "Modeling Government
         Behavior in Collective Bargaining: A Test for Self-Interested
         Bureaucrats." Working Paper #1481

No. 15   Jan K. Brueckner. "Estimating a Bargaining Contract Curve:
         Prior Restrictions and Methodology." Working Paper #1490

No. 16   Peter C. Reiss and Pablo T. Spiller. "Competition and Entry in Small Airline Markets." Working Paper #1497

No. 17   Pablo T. Spiller. "A Note on Pricing of Hub-and-Spoke Networks." Working Paper #1498

No. 18   Larry M. DeBrock. "Joint Marketing Efforts and Pricing Behavior." Working Paper #1500

No. 19   Frank A. Wolak and Charles D. Kolstad. "A Model of Homogenous Imput Demand Under Price Uncertainty."